Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

The Advantages of the UPnP* Internet Gateway Device

Rafael Kolic
Technical Marketing Manager
Corporate Technology Group
Intel Corporation

Prakash Iyer
Senior Architect and Engineering Manager
Intel Labs
Intel Corporation

Overview

Developers who are beginning to learn about Universal Plug and Play (UPnP*) architecture v1.0 should know that there is a lot to be gained by integrating the new technology into their products. Based on open IP standards, UPnP architecture leverages TCP/IP and the Web to enable seamless discovery, control, and data transfer among networked devices. It provides a consistent, interoperable framework for remote Internet Gateway Device (IGD) configuration and management and offers pervasive peer-to-peer network connectivity for PCs of any form factor, for intelligent appliances, and for wireless devices. It can be supported on virtually any operating system and works with virtually any kind of physical networking media, wired or wireless.

Moreover, UPnP architecture is easy to implement. As a result, it is playing a major role in helping to make home and small-business networking more functional and affordable for users. In this article, developers seeking to offer such benefits to their customers can learn about the basics of UPnP architecture and IGD Device Control Protocol (GCP) as well as typical usage scenarios involving UPnP-enabled Internet Gateways.

UPnP* Architecture v1.0

The UPnP Device Architecture is a framework that defines the protocols for communication between controllers, or control points, and devices. UPnP functionality involves five processes:

- Discovery—When a UPnP device is added to the network, the discovery protocol allows the device to advertise its presence to control points by using the Simple Service Discovery Protocol (SSDP). The information exchanged between the device and the control point is limited to discovery messages that provide basic information about the devices and their services (e.g., their types, identifiers, and pointers to more detailed information).
- Description—Using the URL provided in the discovery process, a control point receives XML information about the device, such as manufacturer information like make, model, serial number, and URLs to vendor-specific Web sites. In addition, the description process can also include a list of embedded devices, embedded services, and URLs used to access device features.
- Control—Given knowledge of a device and its services, control points use URLs provided during the description process to access additional XML information that describes actions to which the UPnP device services respond, along with parameters for each action. Control messages are formatted in XML and use Simple Object Access Protocol (SOAP).
- Eventing—When a control point subscribes to a service, the service publishes updates to the control point to announce changes in device status when one or more of the state variables that are evented change. Event messages are formatted in XML and use General Event Notification Architecture (GENA) protocol.
- Presentation—If a UPnP device has an URL for presentation, then the control point can retrieve a page from this URL, load the page into a browser and, depending on the capabilities of the page, allow a user to access interface control features, device, or service information, or any device-specific abilities implemented by the manufacturer.

In addition to the UPnP Architecture, specific UPnP technical working committees, such as the Internet Gateway Device Working Committee, produce additional device-specific specifications for their class of devices, each defined by a Device Control Protocol. The DCP defines variables, actions, and events that allow remote management of such devices. The DCP standardization process includes obtaining three sample implementations of the DCP to pass the UPnP Certification Test Tool, circulating the specification for a mandatory 45-day UPnP Forum member review and comment period, and obtaining the approval of the Steering Committee to become a Standardized DCP.

The UPnP* Internet Gateway Device

The UPnP architecture lends itself well to the discovery, configuration, and management of an IGD. An IGD is an IP-addressable device typically residing at the edge of a home or small-business network. An IGD interconnects at least one LAN with a WAN interface for Internet access. An IGD also provides local addressing and routing services between one or more LAN segments and to and from the Internet.

Intel is supporting IGD functionality through its residential gateways, broadband modems, network adapters, and software solutions. Recently, the company introduced the industry's first gateway product to receive the UPnP certification logo: the Intel® AnyPoint™ Networking Gateway Model 1300. Combining the functionality of a wireless access point, Internet router, and firewall, the AnyPoint Networking Gateway provides seamless discovery and configuration of UPnP-certified applications and devices. The product also extends Intel's industry leadership in UPnP implementation for Linux and UPnP Internet Gateways.

The IGD Device Control Protocol

The IGD DCP (Device Control Protocol) is a set of standardized specifications implemented by UPnP-compliant IGDs. The IGD DCP has a four-pronged focus: Configurable initiation and sharing of Internet connections, advanced connection-management features, management of host configuration services (DHCP), and support for transparent Internet access by non-UPnP-certified devices.

The DCP supports IP and PPP (Point-to-Point Protocol) connections originating at the IGD or transiting the IGD from a host on the LAN; configuration of cable, DSL, POTS, and Ethernet WAN interfaces; remote configuration of LAN DHCP services; and the delivery of both generic and connection-specific status information.

Also included in the IGD specification is a solution for NAT (Network Address Translation) traversal. NAT can "break" many of the compelling new PC and home networking experiences such as multi-player games, real-time communications, and other peer-to-peer services that people increasingly want to use in their homes or small businesses. As IP packets from the private LAN traverse the gateway, NAT translates a private IP address and port number to a public IP address and port number, tracking those translations to keep individual sessions intact.

IGD v1.0: User

Developers considering integrating UPnP architecture in an IGD or other device or application targeting the home or small-business market can think of usage benefits in terms of three functional categories: configuring and querying the state of the IGD, enhancing the ability of applications and devices on the network to access services outside it, and enhancing the ability of services outside the network to access applications and devices on it.

Configuring and querying the state of the IGD:

- Users can access information on the IGD status, the load on any active Internet connection(s), and the available options for installing drivers and/or selected application software without having to remember IP addresses or install custom software.
- Users who are preparing to be away from the network can disable a WAN connection directly from the IGD presentation page and without having to physically disconnect cables.
- Users concerned about line noise or malfunction of their ISP connection can be notified in advance when a connection is in jeopardy of being dropped and when a dropped connection is re-established.
- Users can determine why an Internet connection is malfunctioning, or assist support professionals in doing the same, by accessing connection-debug information that is far more detailed than that traditionally available through operating-system status indicators.

Ensuring that applications and devices on the network can access services outside it:

- Users can launch a UPnP-aware media player that can configure an IGD proxy without requiring user intervention. Similarly, the media player can send IGD bandwidth and other attributes to a media-delivery Web site so the site can select the appropriate presentation stream speed, also without the need for user intervention.
- Users can have their UPnP-enabled Internet telephony application automatically detect and configure the port mapping for Network Address Translation (NAT) on the IGD to enable inbound audio streams, without having to manually configure the IGD.

Ensuring that services outside the network can access applications and devices on it:

- A UPnP-certified video-recorder application can automatically update its program guide nightly with information from a selected Web site by using UPnP to determine the IGD connection status. This capability eliminates the need for such applications to use a private ISP over an analog modem, saving manufacturers and users money and ensuring a higher level of service.
- Users can set up a Web server with UPnP-enabled software that automatically configures the IGD with the necessary port mapping to properly forward all Web requests going through the gateway. UPnP enables the NAT port mapping to take place transparently, without requiring users to manually configure the gateway.

Note the importance of NAT in a few of these usage scenarios. Developers should know that consumers purchasing or leasing an IGD from their ISP are being strongly encouraged to consider only devices that support UPnP for NAT traversal. This capability is essential for lowering support costs and enabling the deployment and use of the most innovative services and applications.

Summary

Incorporating Universal Plug and Play (UPnP) architecture and the new Internet Gateway specification into an Internet gateway device is relatively easy and inexpensive for the device manufacturer. UPnP and the IGD DCP address the problems of NAT traversal and extend the benefit of easier control and accessibility for home and small-business users to almost any application that traverses their network. This improves the user experience for rich broadband applications such as real-time communications, Internet games, and digital entertainment.

More Info

To learn more about UPnP, the IGD DCP, and related specifications, consider joining the UPnP Forum. The UPnP Forum currently has over 400 member companies, many of which are in the process of releasing new UPnP-certified products. More information is available at the Forum's Web site and at the Intel® UPnP Web site.

Author Bios

Rafael Kolic is a technical marketing manager in the Corporate Technology Group's Intel Labs. Since joining Intel in 2000, he has worked on a number of projects involving Universal Plug and Play technology as well as others involving residential gateways, second-screen interactivity, and information management. Rafael has also designed DSP variable speed controllers and performed research in the area of power electronics. Rafael holds an M.E. in electrical and computer engineering and a B.S. in electrical engineering from the University of Florida.

Prakash Iyer is a senior architect and engineering manager in the Network Architecture Lab in Intel Labs. With Intel for over 10 years, he leads a group that develops technologies related to IPv6 routing and transition mechanisms, consumer networking, and seamless, secure IP mobility. Prakash also chairs the Internet Gateway working committee in the Universal Plug and Play Forum. He has previously worked on LAN networking products, IP telephony, and conferencing and VPN technologies. Prakash holds B.S. degrees in physics and electrical engineering and an M.S. in computer science.

Departments

Desktop

Mainstream PC Memory Architectures for the Intel® Pentium® 4 Processor

Scott Jackson
Technical Marketing Engineer
Desktop Product Solutions Division
Intel Corporation

Overview

With the highly anticipated launch of Double-Data-Rate (DDR)-SDRAM support in the Intel® 845 chipset and Intel® Desktop Boards D845BG and D845PT, Intel is offering integrators a way to reach a wider range of price/performance points than ever before—and in so doing will help enable integrators to increase their volumes and reduce their costs. That's because with the introduction of DDR-SDRAM support, integrators now have the option of choosing from the three industry-leading memory architectures: Single-Data-Rate (SDR)-SDRAM for the value segment, RDRAM* for the high-performance segment, and now DDR-SDRAM for the vast mainstream segment in between.

In this article readers will learn the fundamental differences between DDR-SDRAM and SDR-SDRAM and between DDR-SDRAM and RDRAM. Readers also will learn about Column Address Strobe (CAS) latency and its effect on memory performance, and about industry memory naming conventions.

SDR-SDRAM: Classic Value Capabilities

To understand the differences between DDR-SDRAM (commonly known as "DDR") and SDR-SDRAM (commonly known as "SDRAM"), first consider a synopsis of SDRAM, supported in the Intel 845 chipset and Intel Desktop Boards D845HV/D845WN.

A typical PC133 system using SDRAM has a memory bus speed of 133 MHz and a bus width of 8 bytes, providing a theoretical bandwidth of 1,060 MB/sec (133 MHz x 8 bytes), with data sent only on the rising edge of the clock. This bandwidth makes SDRAM ideal for value applications such as office productivity suites, simple Web applications, and value gaming.

Integrators can differentiate their SDRAM systems by considering the different memory timings of memory modules they purchase, specifically the CAS latency. CAS latency is defined as the number of clock cycles from the moment the memory module recognizes a READ command to the availability of the first piece of data. If a READ command is recognized on clock cycle *x*, and the CAS latency is *y*, the data will be available on clock cycle *x+y*; consequently, the lower the CAS latency the higher the performance.

PC133 SDRAM is available in CAS latency 2 and CAS latency 3, for about a $3 (USD) difference in price[1]. Integrators seeking to better exploit the chipset's bandwidth (by as much as 20 percent[2]) can choose SDRAM with CAS latency 2, while those seeking to address an audience more concerned with price can choose SDRAM with CAS latency 3.

[1] Source: Crucial Technology 12/12/01
[2] Source: HotHardware.com

DDR: Mainstream Market Performance

With the launch of DDR support in the Intel 845 chipset, Intel has given integrators a chipset for the vast mainstream performance market. While RDRAM continues to be the high-performance solution, the Intel 845 chipset supporting DDR can hit a wide range of price/performance points within the mainstream market segment. Applications such as office productivity suites, advanced multimedia applications, and mainstream gaming excel on the Intel 845 chipset platform supporting DDR technology.

In many ways the memory architecture of DDR and SDRAM are very similar. To understand the differences, let's examine the available bandwidth of each memory type. This will give a better understanding of the technical and performance differences.

Take, for example, a typical system using DDR266, similar to SDRAM. The memory bus width is 8 bytes and the memory clock frequency is 133 MHz. While an SDRAM module running with the aforementioned specifications would provide a bandwidth of 1,060 MB/s, a DDR module, due to its ability to transfer 2x the data per clock cycle, will double the bandwidth to 2,100 MB/s (133 MHz x 8 bytes x 2 transfers/cycle). This increase in bandwidth allows the chipset to better fill the Pentium® 4 processors system bus, resulting in increased performance on many applications.

A less significant difference between DDR and SDRAM involves CAS latency, which is available in latencies of 2 and 3 for SDRAM, and 2 and 2.5 for DDR. As with SDRAM, a lower CAS latency will result in an increase in performance but at additional cost.

RDRAM: Still the Top Performer

Even with Intel's introduction of support for DDR, integrators working in the high end will continue to design their systems around RDRAM, supported in the Intel® 850 chipset and Intel Desktop Boards D850MD and D850MV. Applications such as 3D modeling, video editing/encoding, audio encoding/decoding, and high-end gaming take advantage of RDRAM's unmatched memory bandwidth and performance.

To understand why RDRAM is still the right choice for the high-performance market segment, let's examine the potential bandwidth an Intel 850 chipset-based system can provide. A single RDRAM memory channel operates at 400 MHz with a bus width of 2 bytes and, very much like DDR technology, RDRAM can transfer data on the rising and falling edges of a clock cycle. Therefore, memory bandwidth for a single channel of RDRAM is 400 MHz x 2 data read/cycle x 2 bytes bus width, or 1,600 MB/s.

The Intel 850 chipset is unique in that it can control two independent memory channels; this is why an integrator needs to buy RIMMS in matched pairs and populate both channels. Since there are two memory channels acting independently, this increases the bandwidth from 1,600 MB/s/channel x 2 channels to offer a bandwidth of 3,200 MB/s. This memory bandwidth provides a perfect balance to the Pentium 4 processor's system bus bandwidth, which is also 3,200 MB/s. Balance is one of the reasons why RDRAM is still positioned as the high-performance platform for Pentium 4 processor-based systems and why it still provides best-in-class performance for today and tomorrow's performance desktops.

A Note about Memory Naming

Whether integrators are shopping for SDRAM, DDR, or RDRAM for their next systems, they need to understand memory naming conventions, specifically how the naming followed by Intel and several other major entities differs from methods followed by other organizations.

In naming DDR-based systems, Intel follows the same convention used for SDRAM-based systems, with a slight modification. For standard SDRAM, the value of *x* in the term "PC*x*" refers to the system's memory-bus speed, in megahertz, times the number of reads per clock cycle. Intel endorses the use of this nomenclature for DDR except that "DDR" replaces the "PC" as the prefix to the rating. For example, DDR200 and DDR266 are for memory modules running at effective bus speed of 200 and 266 MHz respectively.

In a differing naming convention being used for DDR, the *x* in "PC*x*" refers to the system's effective bandwidth in megabytes. This means that the DDR-based systems described above might not be labeled DDR200 and DDR266 but rather PC1600 and PC2100.

Lastly, although the memory modules for DDR are similar to the memory modules for SDRAM, they have different pin counts and are keyed differently. This means that existing boards designed for SDRAM (boards that integrators may have in stock, for example) will not support DDR memory modules.

Summary
With its introduction of the Intel 845 chipset and Intel desktop boards supporting DDR-SDRAM in addition to SDR-SDRAM, Intel is taking the next step in its technology and manufacturing leadership. Added to the support for RDRAM provided by the Intel 850 chipset and Intel desktop boards, this means that no matter what market or application a given integrator needs to target, Intel has the right chipset and motherboard to satisfy that need.

More Info
For motherboard information, visit the Intel® Desktop Boards area of the Intel® Developer Site.
For chipset information, visit the Intel® Chipsets area of the Intel Developer Site.

Author Bio
Scott Jackson is currently a member of the Intel Sales and Marketing Rotation Program. His current assignment is with the Desktop Platform Solutions Division as a product technical marketing engineer. Previously, Scott has held product marketing engineer positions in the same division and in the Cellular Communication Division.

Before joining Intel last year, Scott worked for Digital Equipment Corporation and Compaq Computer Corporation as a validation engineer on the Alpha processor family. He holds a B.S.E. in computer engineering from the University of Michigan.

Active Hardware Monitoring and Control for Quieter Desktop PCs

Mike Castillo
Desktop HW Manageability Manager
Desktop Platform Solutions Division
Intel Corporation

Syed Ahmed
Strategic Marketing Engineer
Desktop Platform Solutions Division
Intel Corporation

Overview

The need for quieter PCs is quickly increasing due to consumer demands and existing PC specifications. Government specifications, such as those from Japan, Germany, and Sweden, now include significant acoustic requirements. Industry specifications, such as the Blue Angel specification and TCO'99, also place increasing demands on developers to lower PC acoustic levels. Yet these specifications are increasingly difficult to meet since systems with high-performance processors consume increasing amounts of power, and thus dissipate tremendous amounts of heat. Simply managing system thermals by increasing the fan solution is no longer good enough. Designing an acoustically mindful, thermally efficient PC for today and tomorrow requires active hardware fan control.

To comply with consumer demands and government requirements, Intel has developed an ASIC specification to help designers meet PC acoustic requirements. The specification defines an ASIC that not only monitors but also controls system fans to reduce their impact on system acoustics. The ASIC also monitors temperatures, voltages, and fan tachometers.

Two ASIC manufacturers are already offering this chip to developers. A third manufacturer expects to have products available in early 2002. With the launch of Intel's new 845 family chipsets, as well as higher Pentium® 4 processor speeds, the new ASIC is welcome and needed.

Fan Control

In the past, developers have dealt with growing thermal requirements by increasing the number of system fans and by increasing the size of the processor heat sink. However, without intelligent control, the fans must continuously run at full speed, even when the system is in an idle state.

Under normal conditions, a typical user may not run applications that exercise the processor, system, and memory so rigorously that the fans need to be at full speed. Word processing, Web surfing, and e-mail, for example, generate lower levels of power and should not require that fans operate at full speed. Other applications, such as gaming and DVD movies, are processor- and memory-intensive, and they require fans operating at higher speeds to provide proper cooling.

The less intensely the processor and memory are used, the lower the cooling requirements. Less cooling means a slower fan speed. In some situations, such as an idle or suspend-to-RAM state, the chassis fan—and possibly all fans—may be turned off. Hardware-based control of the fans lets a system react to actual usage, so that it's cooled only as needed.

Why Hardware?

Previously, developers have used software-based applications to control fans. As with hardware management, these applications used measurements of ambient air temperature or measurements of CPU temperature to determine fan speed.

However, implementing fan control through software is problematic. If the PC, operating system, or fan-control application hangs, the result could be an unexpected hardware failure.

In a software-hang scenario, the fan speed would stay at the last programmed value prior to the hang, regardless of the system's cooling requirements. Thus, the system could overheat, causing components to be damaged or destroyed. Fan speed control should always be implemented through hardware, rather than software.

Profiling

Regardless whether the fan control solution is implemented in hardware or software, system thermals must be properly characterized before effective fan-control can occur. This means accurately answering questions such as: How much heat does your graphics solution generate? At what measured temperature should each fan spin-up? What should be the fan-speed ramp rate for coming up to full speed? And at what temperature should all fans be brought to full speed?

Other system- and board-level details should also be considered. If the length of the traces between the diode and ASIC are long, the reported temperature is reduced due to the increased resistance of the trace. If these traces are long or close to noisy signals, the temperature that was “read” can be significantly inaccurate. To define the most acoustically effective solution, thermal characterization of the system should be done in a thermal-test chamber to determine accurate temperatures when the fans need to turn on or off.

The New ASIC

The new ASIC is a hardware-monitoring and fan-control ASIC for PC fans. The new ASIC not only monitors the tachometer, but controls processor and chassis fans.

Thermal sensors are placed in strategic areas on the motherboard where the fans can affect temperature. For example, the ASIC can use the processor’s on-board thermal diode and a thermal diode placed on the motherboard to capture the ambient temperature inside the chassis. If the system has other hot spots—such as a powerful graphics or memory solution—thermal diodes can be placed in those "temperature zones" also. The ASIC uses measurements from those diodes to throttle the speed of the fans that affect those zones.

One of the main features of the new ASIC is that it can control each of these zones separately. For example, if the graphics solution is generating more heat, the fan that affects this temperature zone can be accelerated without also speeding up other fans. During idle periods, the system fans may be decelerated, while only the processor fan remains on at a low speed. By controlling the fan for each temperature zone individually and as needed, developers can create a PC with the quietest possible environment.

ASIC Features

The new ASIC has great features:

- Two thermal diode inputs (a CPU diode input and a miscellaneous system diode input) plus an on-board ambient temperature sensor, providing a total of three thermal input zones
- Control of up to four fans through three pulse-width modulation (PWM) outputs
- Separate control of the three temperature zones
- Four fan tachometer inputs
- An absolute temperature limit so that all fans turn on at full speed if any temperature zone (including the processor) rises above a critical limit
- Temperature hysteresis, with a range of up to 15°C, to prevent fans from toggling on and off at temperature boundaries
- The ability to turn the chassis fans off completely
- A smoothing feature to compensate for the CPU's tendency toward temperature spikes during burst usage
- Different PWM frequencies to adjust for a specific fan's characteristics
- Register lock, to prevent viruses or rogue code from changing the defined settings
- Status and error messaging
- Backward compatibility with previous fan-management ASICs (that conformed to Intel specifications)

Developers can implement up to three fan-control ASICs on the SMBus (system management bus) at any time to monitor additional processors and temperature zones.

PWM Features

One of the most important features on the new ASIC is the ability to control up to four fans using its three PWM (pulse width modulated) outputs. The outputs are designed to control the CPU fan, a secondary fan (for the memory subsystem as needed or for another system component), and one or two chassis fans. When there are two chassis fans, both fans are controlled by a single PWM output.

The new ASIC also lets developers assign more than one PWM output to a single thermal zone. For example, PWM outputs 2 and 3, connected to two chassis fans, may both be controlled by the temperature readings from the remote thermal diode. If the ambient temperature at the remote thermal diode rises to a certain level, the fans cooling zones 2 and 3 could be throttled up without affecting the fan cooling another zone. If the ambient temperature near the remote thermal diode exceeds the absolute limit, all PWM outputs would then go to 100 percent duty cycle to provide maximum cooling for the system. Additionally, a fan may be controlled by the hottest of several zones. For example, the hottest of the on-board thermal sensor or remote diode may be used to control the chassis fans.

Developers also have the option of using different PWM frequencies, from 10 Hz to 94 Hz, for the PWM outputs. This is useful for adjusting system acoustics for a specific fan, since some fans may be quieter if pulsed faster or slower.

Spin-up Times

Spin-up times should also be considered when designing fan control solutions. Usually, when a fan is off, there is a turn-on spike in which the fan is brought to full power to overcome inertia. Full-speed continues for a specific period of time, such as for 350 ms, until it is assumed that the fan is spinning. The fan speed is then decreased to the specifically required level.

For three-wire fans, the new ASIC does not turn the fan on full-speed for a specific amount of time. Instead, the new ASIC uses tachometer inputs to determine fan speeds. The PWM outputs are decreased as soon as the appropriate speed is reached, usually within a much shorter period than a safe default spin-up time.

Summary

Domestic, international, and industry standards are all demanding better PC acoustics. With new Pentium 4 processor platforms and with Intel's new 845D chipset pushing PC performance, it is becoming increasingly important for PC system designers to create a balance between thermals and acoustics. In view of demands for quieter PCs, developers should begin designing their systems to include the new fan-control ASIC.

Today's performance system is tomorrow's mainstream PC. Intel® Desktop Boards offer hardware-based fan-control today to create quiet, thermally efficient platforms that can handle tomorrow’s power-dissipation demands. Active control, not just monitoring of system thermal levels, will help developers reduce overall system acoustic levels.

More Info

There are several industry specifications for PC acoustics. The Blue Angel industry specification and the Japanese government specification for acoustic levels can also be found online.

For more information visit the Motherboard section of the Intel® Developer Site.

Fan-control ASICs that comply with the new specification are currently available from these companies:
Analog Devices, Inc. (ADI)
Part Number: ADM1027ARQ
ADM1027 datasheet
General temp sensor selector guides

Standard Microsystems Corporation (SMSC)
Part Number: EMC61D101-CK
EMC61D101 datasheet

National Semiconductor
LM85CIMQX (tape & reel)
General temp sensor selector guides

Disclaimer: Implementation of the technology set forth in this article may require licenses from third parties.

Author Bios

Mike Castillo has worked for Intel for six years, first as a design engineer, then as a desktop hardware manageability manager. His team developed the Intel specification for the new hardware monitoring and fan control ASIC. Mike has worked on previous ASICs in this family as well as the Port Angeles project, which integrates Super I/O (SIO), glue logic, and manageability functions. Mike has received several divisional awards for his work on PC100 enabling, modem-on-motherboard, and HPNA-on-motherboard technologies. He holds two patents for flash memory. Mike received his B.S. from Brigham Young.

Syed Ahmed has worked at Intel for five years and currently is a platform marketing engineer in the Desktop Platform Solutions Division strategic marketing group. Syed helped develop the Intel specification for the new hardware management and control ASIC for PCs. He worked on the Integrated SIO and Glue projects and created design guide for standardized front-panel I/O connectivity for internal and external PC devices. Syed has taught classes at Intel field training events and at the Intel Developer Forum Conference and he's earned several divisional awards for work on the PC100 and various time-to-market platforms. He received his B.S.E.E. from The Wichita State University, and his M.B.A. from the University of Phoenix.

—End of Intel Developer Update Magazine Issue 28—